

FIGURE 1

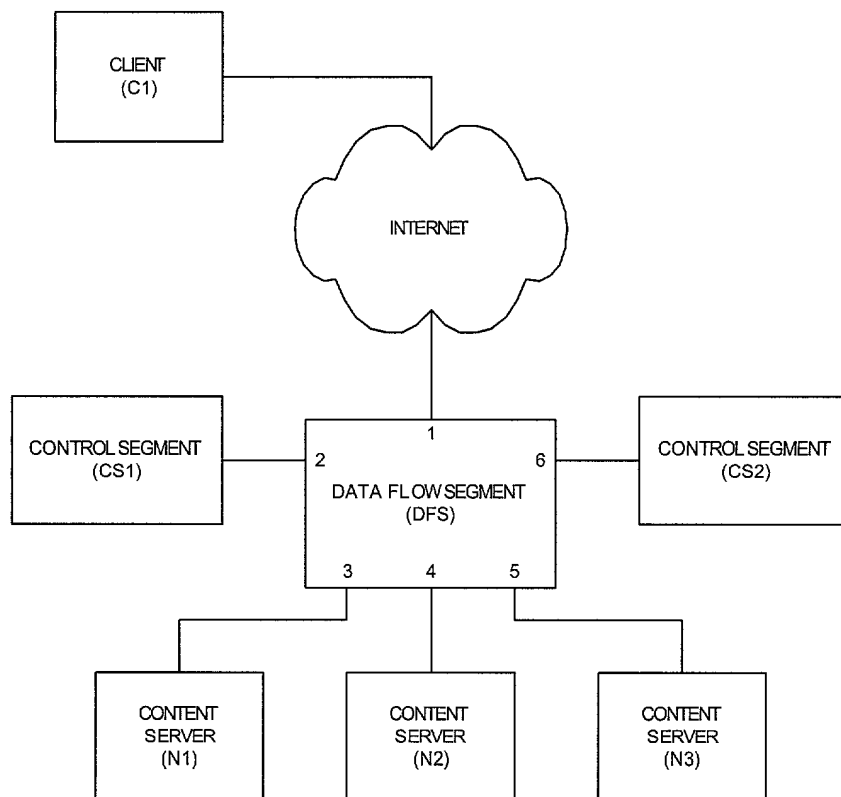


FIGURE 2

TABLE I : Message Types

Mnemonic	DFS Sends?	CS Sends?	Priority	Batching?	Single mbuf?	State Sharing?	Expected Response	Explanation
QUERY	Y	N	H	N	Y	N	REPLY	Request from the DFS to the CS for instructions on handling a new flow.
APP_QUERY	Y	N	H	N	N	N	REPLY	QUERY message for a virtual server that is load balanced using application data.
REPLY	N	Y	H	N	Y	Y	-	Instructions provided by the CS to the DFS on handling the new flow.
DELETE	Y	N	H	N	Y	Y	-	Notification from the DFS to the CS that a flow was deleted due to inactivity, expired timer, TCP connection close or reset.
REAP	N	Y	L	Y	N	N	STATS	Notification from the CS to the DFS to delete a flow due to expired timer.
STATS	Y	N	L	Y	Y	Y	-	Statistics sent from the DFS to the CS, when the CS has requested a flow to be deleted.
NEWFLOW	N	Y	L	N	N	Y	-	Notification to the DFS from the CS to prepare to handle a new flow.
RESET	Y	Y	L	N	N	Y	ACK (CSMB)	Delete all flows.

FIGURE 4

TABLE II: Message Fields

Field	Length (octets)	Field Name
Message Type (see Fig. 4)	1	msg_type
Global Serial Number	2	serial_global
Serial Number in this type	2	serial_bytype
Original Source IP address	4	orig_src_ipaddr
Original Source TCP or UDP port	2	orig_src_port
Original Destination IP address	4	orig_dst_ipaddr
Original Destination TCP or UDP port	2	orig_dst_port
New Destination IP address	4	new_dst_ipaddr
New Destination TCP or UDP port	2	new_dst_port
Next Hop IP address	4	next_hop_ipaddr
TCP Server Sequence Number Offset	4	svr_seq_offset
TCP Server Acknowledgement Number Offset	4	svr_ack_offset
TCP Client Sequence Number Offset	4	client_seq_offset
TCP Client Acknowledgement Number Offset	4	client_ack_offset
Flags (see Fig.	1	flags
TCP Flags	1	tcp_flags
Type of Service	1	tos
Protocol	1	proto
Error code	1	error
Application data length	2	app_data_len
Application-specific data	(varies)	app_data
DFS Vendor ID	4	vendor_id
Free test string	(varies)	-

FIGURE 5

TABLE III: Message Contents

Message Type	Fields											
	app_data_len	app_data	error	proto	tos	tcp_flags	flags	svr_seq_offset, svr_ack_offset, client_seq_offset	client_ack_offset	next_hop_ipaddr	new_dst_port	new_dst_ipaddr
QUERY	N	N	N	Y	Y	Y	N	N		N	N	N
APP_QUERY	Y	Y	N	Y	Y	Y	N	N		N	N	N
REPLY	N	N	Y	Y	Y	Y	Y	Y		Y	Y	Y
DELETE	N	N	N	Y	Y	Y	Y	Y		Y	Y	Y
REAP	N	N	N	Y	Y	Y	Y	Y		Y	Y	Y
STATS	N	N	Y	Y	Y	Y	Y	Y		Y	Y	Y
NEWFLOW	N	N	N	Y	Y	Y	Y	Y		Y	Y	Y
RESET	N	N	N	N	N	N	N	N		N	N	N

FIGURE 6

Flags field boolean variables

Variable	Explanation
OUTBOUND	If TRUE, then the message is regarding an outbound flow. If FALSE then the message is regarding an inbound flow. Most messages are inbound flows.
ADD_TCP_OFFSET	If TRUE, then the values of the seq_offset and ack_offset fields should be arithmetically added to the appropriate values in the TCP header of every packet in the flow. This is typically used when the TCP handshake proxy is performed by the CS. If FALSE, then those fields should be ignored.

FIGURE 7

Error codes for the error field

Value	Explanation
UNKNOWN	The CS does not know how to handle the flow in the query.
NOTAVAIL	CS REPLY to QUERY message. The virtual server is not available due to port denied or maintenance mode. The flow is denied.
CONNLIMIT	CS REPLY to QUERY message. A user defined connection limit for a virtual server would have been exceeded. The flow is denied.
CONNALIVE	DFS STATS reply to SSMB REAP message. The connection is still alive and should not be reaped on the CS. Statistics should be incremented on the CS.

FIGURE 8

SSMB Message Header

0	1	2	3	4	5	6	7
flow serial (4)				msg_type	court	error_code (2)	

Flow Messages

0	1	2	3	4	5	6	7
SSMB Header(8)							
orig_src_ipaddr(4)				orig_dst_ipaddr(4)			
orig_src_port(2)		orig_dst_port(2)		tos	prdb	flags	tcp_flags
new_dst_ipaddr(4)				new_dst_port(2)		pad(2)	
next_hop_ipaddr(4)				svr_seq_offset(4)			
svr_ack_offset(4)				client_seq_offset(4)			
client_ack_offset(4)				flow_id(4)			

Application Proxy Query Message

0	1	2	3	4	5	6	7
SSMB Header(8)							
orig_src_ipaddr(4)				orig_dst_ipaddr(4)			
orig_src_port(2)		orig_dst_port(2)		tos	prdb	flags	tcp_flags
new_dst_ipaddr(4)				new_dst_port(2)		pad(2)	
next_hop_ipaddr(4)				flow_id(4)			
pad(2)		app_data_len(2)		app_data(app_data_len)			

Endflow Messages

0	1	2	3	4	5	6	7
SSMB Header(8)							
orig_src_ipaddr(4)				orig_dst_ipaddr(4)			
orig_src_port(2)		orig_dst_port(2)		tos	prdb	flags	tcp_flags
inbound_bytes(8)							
outbound_bytes(8)							
inbound_pkts(8)							
outbound_pkts(8)							
flow_id(4)							

FIGURE 9

TABLE IV : Message Types

Mnemonic	DFS Sends?	CS Sends?	Expected Response	Explanation
ADD_VS	N	Y	-	Add or modify a virtual server.
DELETE_VS	N	Y	-	Delete virtual server. If the virtual server does not exist, the message is ignored.
QUERY_VS	Y	Y	VS_LIST	Send list of all virtual servers.
VS_LIST	Y	Y	-	List of currently configured virtual servers.
RESET	Y	Y	-	Delete all virtual servers. Typically sent by the CS when it initializes.
STATUS	Y	Y	STATUS	Miscellaneous house keeping functions.
HELLO	Y	Y	ACK	Miscellaneous house keeping functions.
ACK	Y	Y	-	Miscellaneous house keeping functions.

FIGURE 10

TABLE V : Message Fields

Field	Length (octets)	Field Name
Message Type (see Fig. 10)	1	msg_type
Global Serial Number	2	serial_global
Serial Number in this type	2	serial_bytype
Message Length (in bytes, including payload data but not TCP or lower-layer headers)	4	msg_length
Major Version Number	1	vers_major
Minor Version Number	1	vers_minor
Virtual IP address	4	virt_ipaddr
Virtual TCP or UDP port	2	virt_port
Protocol	1	proto
Virtual Server Class	1	vs_class
Flags	4	flags
Error Code	1	error
SSMB IP address for active CS	4	ssmb_active_ipaddr
SSMB IP address for standby CS	4	ssmb_standby_ipaddr
SSMB UDP port	2	ssmb_port
Application configuration data (HTTP cookie name, etc.)	(varies)	app_data
Application specific data length in bytes	2	app_data_len
Free text string	(varies)	-

FIGURE 11

TABLE VI : Message Contents

Message Type	Fields													
	auth_data_len	auth_data	smb_port	smb_standby_ipaddr	smb_active_ipaddr	app_data_len	app_data	proto	error	flags	vs_class	protp	virt_port	virt_ipaddr
ADD_VS	N	N	N	N	N	Y	Y	Y	N	Y	Y	Y	Y	Y
DELETE_VS	N	N	N	N	N	N	N	N	N	N	N	N	Y	Y
QUERY_VS	N	N	N	N	N	N	N	N	N	N	N	Y	Y	Y
VS_LIST	N	N	N	N	N	Y	Y	Y	N	Y	Y	Y	Y	Y
RESET	N	N	N	N	N	N	N	N	N	N	N	N	Y	N
STATUS	-	-	-	-	-	-	-	-	-	-	-	-	-	-
HELLO	Y	N	Y	Y	Y	N	N	N	N	Y	N	N	N	N
ACK	-	-	-	-	-	-	-	-	-	-	-	-	-	-

FIGURE 12

Virtual Server Classes

Value	Explanation
CS_ASSIST	A CS-assisted virtual server.
DFS_ASSIST	A DFS-assisted virtual server.

FIGURE 13

CSMB Header

0	1	2	3	4	5	6	7
global_serial(2)		msg_type_serial(2)		msg_length(4)			
msg_type	pad	error_code(2)					

Virtual Server Definition

0	1	2	3	4	5	6	7
flags (4)				virt_ipaddr (4)			
virt_port(2)		vs_class	proto	pad (2)		app_data_len(2)	
app_data (app_data_len)							

Hello Message

0	1	2	3	4	5	6	7
CSMB Header (12)							
				ssmb_active_ipaddr(4)			
ssmb_standby_ipaddr(4)				flags (4)			
vers_majr	vers_minr	ssmb_port(2)					

Virtual Server Definition

0	1	2	3	4	5	6	7
flags (4)				virt_ipaddr (4)			
virt_port(2)		vs_class	proto	pad (2)		app_data_len(2)	
app_data (app_data_len)							

Status Messages

0	1	2	3	4	5	6	7
CSMB Header (12)							
				status (4)			

FIGURE 14

Flags field boolean variables

Variable	Explanation
TRANSLATE_ADDR	If TRUE, then the CS will always supply a new destination IP address for each flow. If FALSE, then the DFS should not translate source or destination IP addresses for any flow on this virtual server, and the DFS should always ignore the new_dst_ipaddr field of the REPLY message from the CS.
TRANSLATE_PORT	If TRUE, then the CS will always supply a new destination TCP or UDP port for each flow. If FALSE, then the DFS should not translate source or destination ports for any flow on this virtual server, and the DFS should always ignore the new_dst_port field of the REPLY message from the CS.
ROUTE_BY_DST_IP	If TRUE, then the DFS will perform a route lookup on the destination IP address to determine the next hop route for the flow. If FALSE, then the DFS will use the next_hop_ipaddr field to determine the next hop for the flow.
REDUNDANT	(HELLO message) If TRUE, then there are two CSs and the ssmb_standby_ipaddr will be used for state sharing. If FALSE, then either the CS is not redundant or state sharing is not desired.
WILDCARD_ADDR	If TRUE, then the virt_ipaddr field is ignored and all traffic received from the external network destined for any address which does not match another virtual server or other known local address is processed as if it was addressed to this virtual server.
WILDCARD_PORT	If TRUE, then the virt_port field is ignored and all traffic destined for this virtual address that does not match another virtual port of any other virtual server is processed as if it was addressed to this virtual server.
NOARP_MODE	If TRUE, then the controller acts like a router and accepts packets destined to this address but does not respond to ARP requests for it. If FALSE, then the controller acts as a host and advertises the address (i.e. response to ARP requests).
APP_PROXY	(HELLO message) If TRUE, then the controller supports application data load balancing and can perform the TCP handshake proxy as well as extract application data from the client request. The CS always sets this flag to true. The DFS sets this flag to true if it has sufficient capability.
SSL_PROXY	If TRUE, the CS makes load balancing decisions for the virtual server based upon the client's SSL session id and sends it to the CS.
COOKIE_PROXY	If TRUE, the CS makes load balancing decisions for the virtual server based upon the value of a cookie in the HTTP request. The DFS should proxy the client connection, extract the designated cookie, and send the cookie to the CS. The cookie name is provided in the app_data field.
HTTP_PROXY	If TRUE, the CS makes load balancing decisions for the virtual server based upon the value of the HTTP request. The DFS should proxy the client connection, extract the HTTP request, and send the data to the CS.
RAW_PROXY	If TRUE, the CS makes load balancing decision for the virtual server based upon the application data format that is not supported by the DFS. The DFS should proxy the client connection, and bridge packets that are received from the client to the CS.

FIGURE 15

Error codes for the error field

Value	Explanation
VERS_NEW	The version specified in the HELLO message is not yet supported by this segment. The other segment should send another HELLO message with a lower version number, if possible.
VERS_OBSOLETE	The version specified in the HELLO message is no longer supported by this segment. The other segment should send another HELLO message with a higher version number, if possible.

FIGURE 16

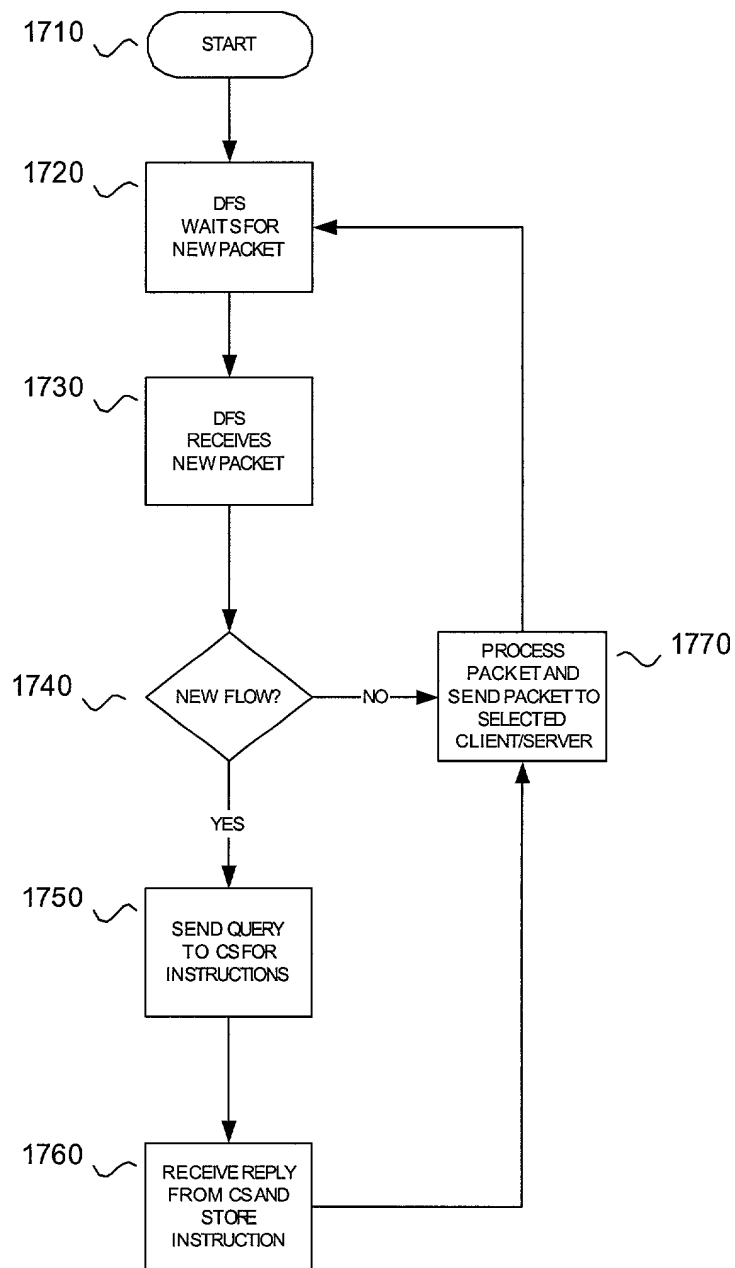


FIGURE 17

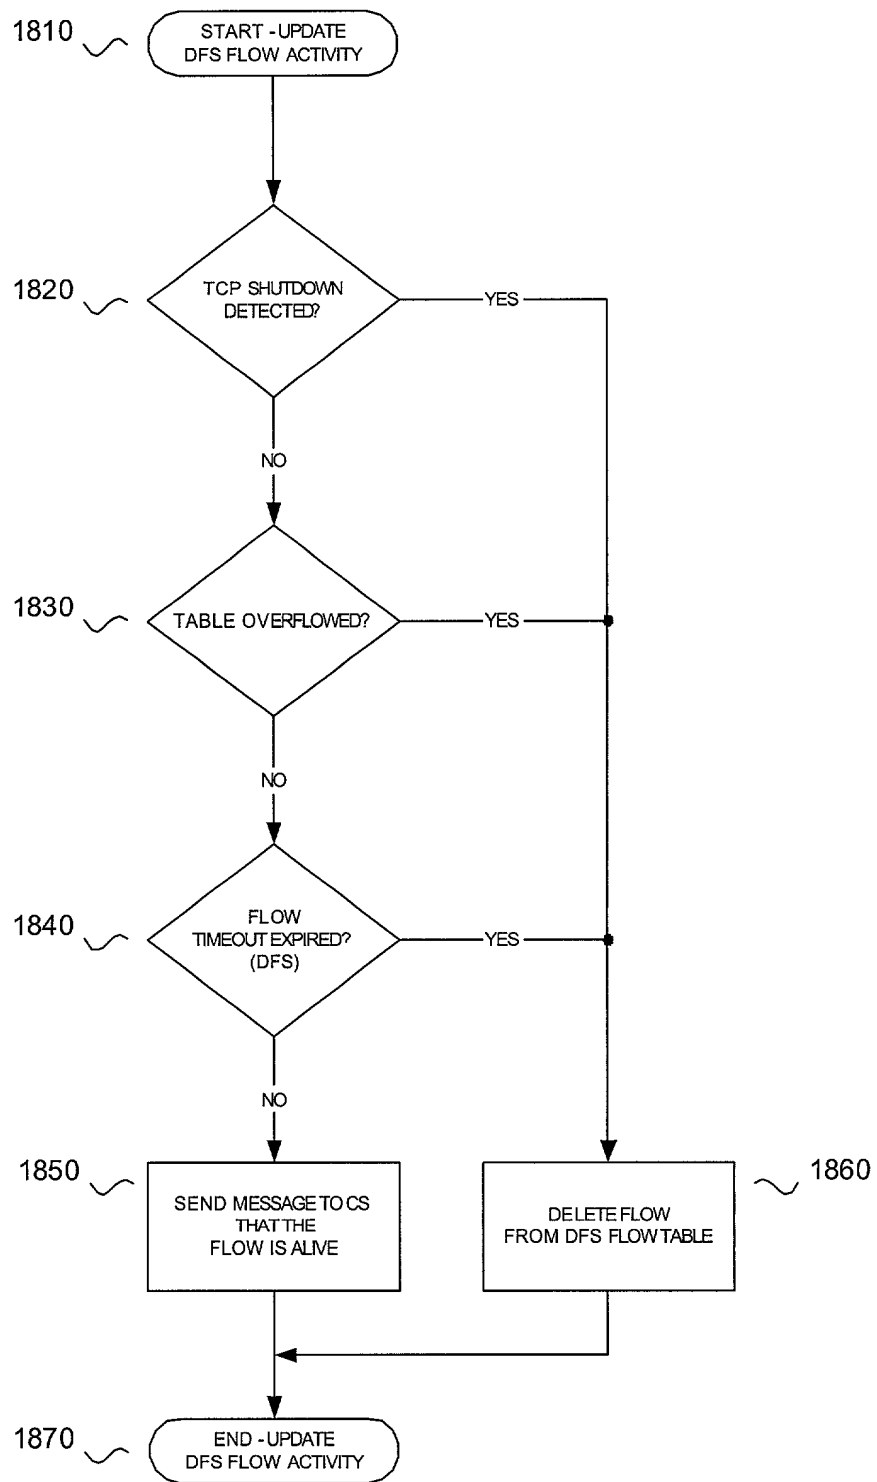


FIGURE 18

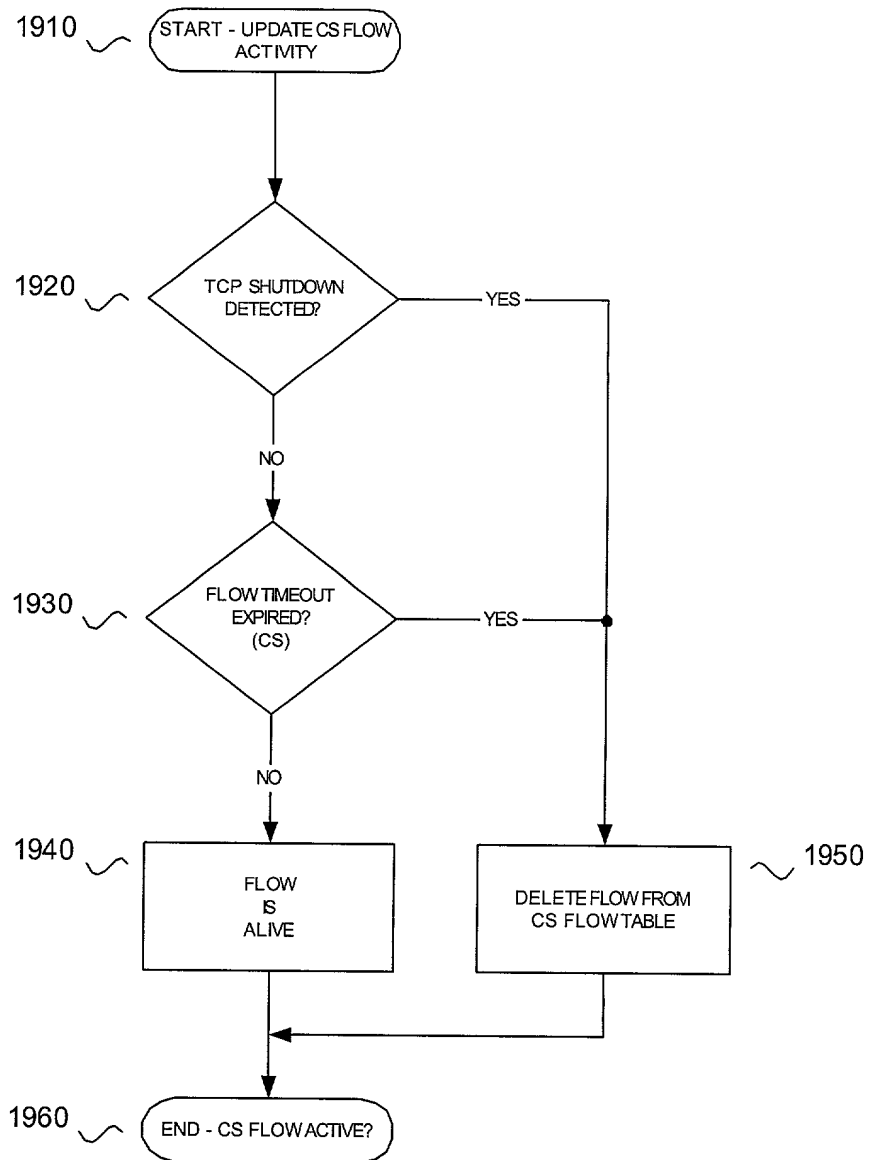


FIGURE 19

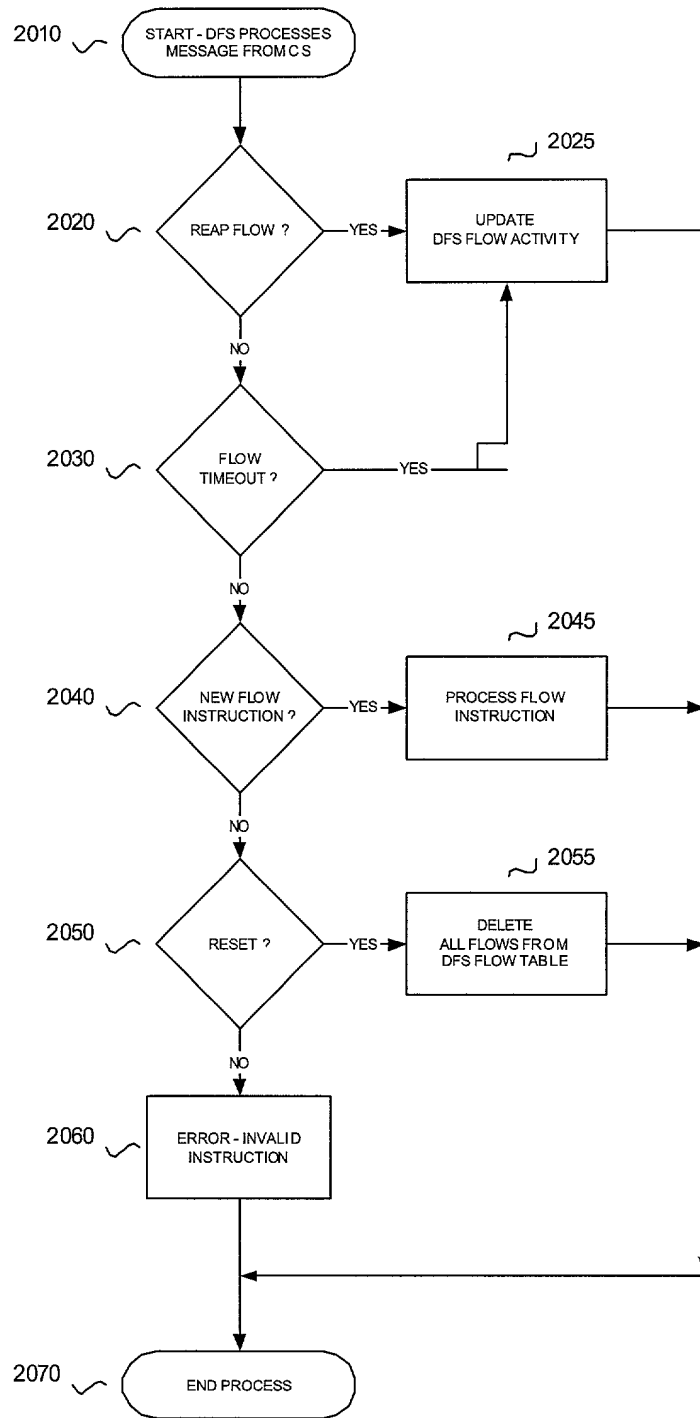


FIGURE 20

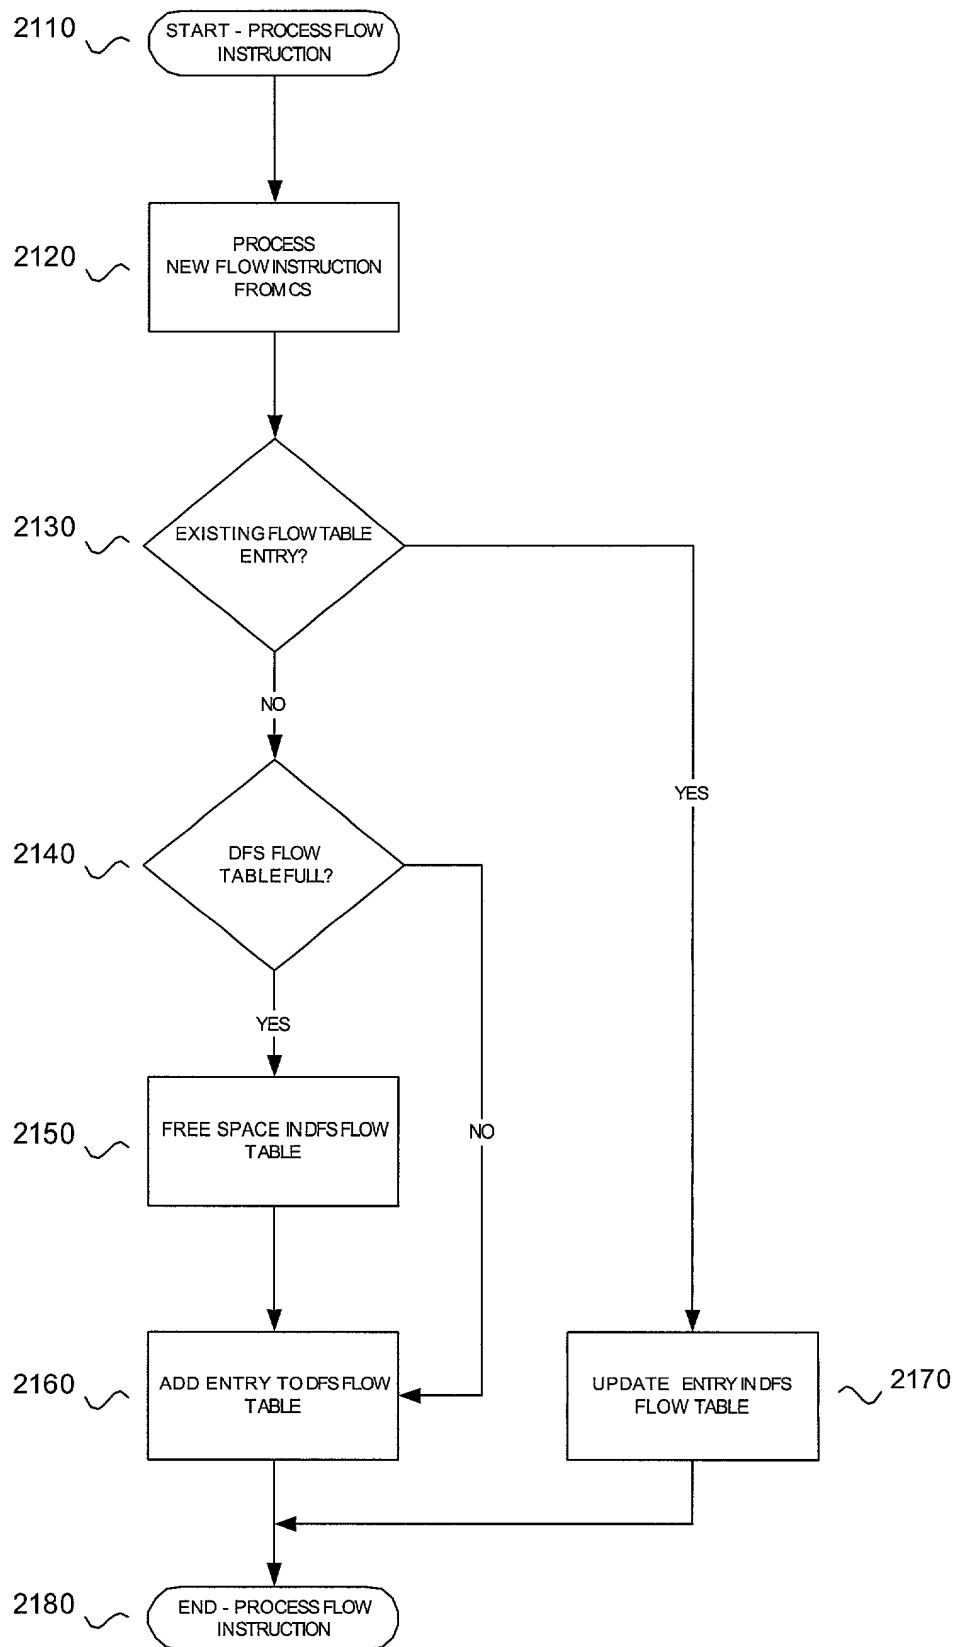


FIGURE 21

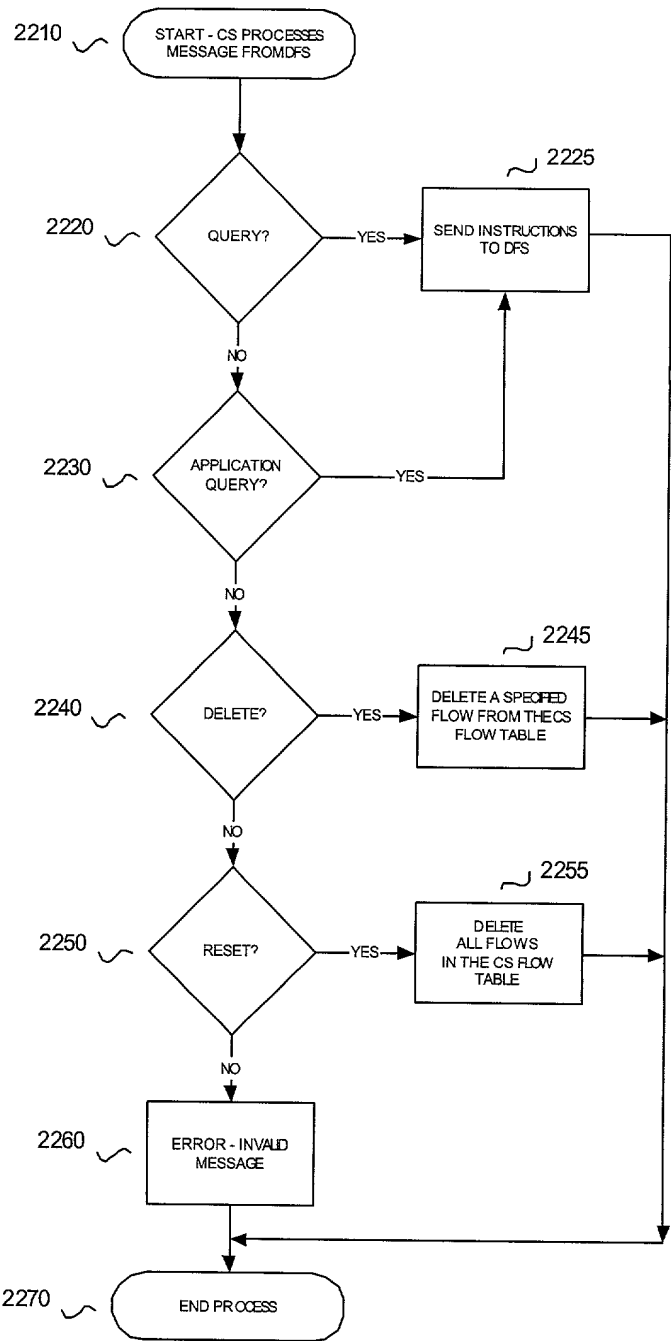


FIGURE 22